McKinsey&Company

McKinsey Working Papers on Risk, Number 49

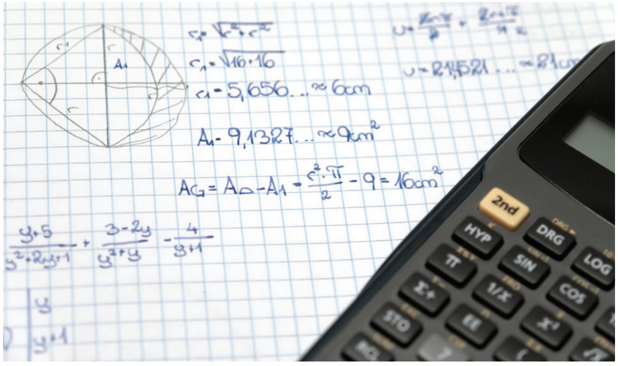

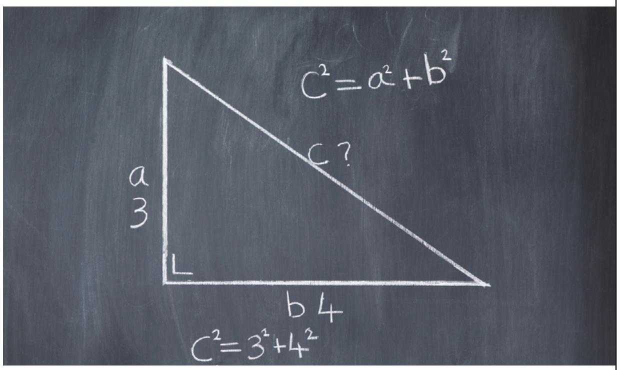

The hypotenuse and corporate risk modeling

Martin Pergler

August 2013

© Copyright 2013 McKinsey & Company

Contents

McKinsey Working Papers on Risk presents McKinsey's best current thinking on risk and risk management. The papers represent a broad range of views, both sector-specific and cross-cutting, and are intended to encourage discussion internally and externally. Working papers may be republished through other internal or external channels. Please address correspondence to the managing editor, Rob McNish (rob_mcnish@mckinsey.com).

The hypotenuse and corporate risk modeling

"I'm teeming with a lot o' news
With many cheerful facts about the square of the hypotenuse."

—*W.S. Gilbert*[1]

"…buys movie rights for six million rubles, changing title to *The Eternal Triangle*, with Ingrid Bergman playing part of hypotenuse."

—*Tom Lehrer*[2]

If you're a risk manager, you probably rarely think about the hypotenuse. Maybe you remember from school days that it's the long side of a right-angled triangle. Maybe the word conjures up one of the songs above. But if you're building probabilistic models for your company, the hypotenuse can save you time—and maybe even save your company money.

Increasingly, real-economy companies (corporates) are modeling the volatility of key metrics in their business plan: cash flow, earnings, portfolio value, and a range of other financial metrics. The results can be quite powerful. What is your company's likelihood of running out of money, being downgraded, or being able to meet growth aspirations? And what can be done to shape it? Hedging, contracting changes, adjusting the financial structure, prioritizing the growth portfolio—all of these can be looked at in an integrated fashion.

The drivers of these models are external variables, such as currency exchange rates and commodity prices, and internal ones, such as operational efficiency, uptime, yield, and project costs or delays. Companies that start building these models quickly come up with 20 to 25 such drivers, even ignoring harder-to-quantify strategic or operational risks.

The reality is that typically only a handful of these drivers really matter and are worthwhile to model in detail. The reason is the hypotenuse.

The Major-General's most important fact

Gilbert's Major-General (perhaps fortunately) does not share his cheerful facts about the hypotenuse. For us, the most important cheerful fact is the following: the length of the hypotenuse is disproportionally affected by the length of the triangle's longer side. Consider Exhibit 1.

Exhibit 1 The triangle's longer side disproportionally affects the length of the hypotenuse.

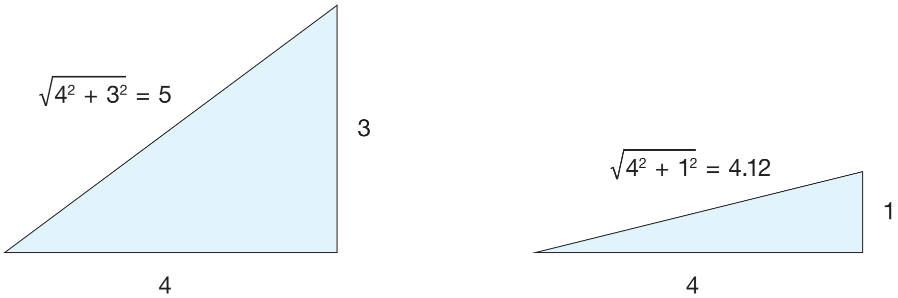

1 "I Am the Very Model of a Modern Major-General," *The Pirates of Penzance*.
2 "Lobachevsky," *Songs by Tom Lehrer*.

The length of the hypotenuse of each triangle follows the law of the sum of squares (the Pythagorean theorem). That means that in the triangle at left, the 3 side, just a bit shorter than the 4 side, only contributes 36 percent (3^2 divided by 5^2) of the length of the hypotenuse. If the 3 drops to 2, the hypotenuse shortens to 4.47; if the 4 drops to 3, it shortens more, to 4.24 (all calculated with the same square-root and sum-of-squares formula). So the longer side is more important. The difference is even more extreme in the skinnier triangle at right: the 1 side contributes very little to the hypotenuse (4.12)—its length basically all comes from the 4 side.

What does this have to do with risk modeling? Suppose your business had only two volatility drivers, both normally distributed and completely independent (strong assumptions, I agree). Probability theory implies that the combined level of volatility—to your cash flows or other financial metrics—follows the same law of the sum of squares, just like the hypotenuse. If the first driver contributes $3 million of risk to your bottom line (whether that's a "standard deviation" measure or something like a 95th percentile value at risk) and the second contributes $4 million, the two together will contribute $5 million. At the heart of this is the somewhat esoteric fact that the so-called variance of the sum of two independent normally distributed random variables is the sum of the two variances. The result is that, just as we saw for the triangles, the $3 million risk contributes a bit less than its fair share than the $4 million one; if it is only $1 million instead of $3 million, it contributes next to nothing.

Blasting the Major-General into hyperspace to save modeling time

Now suppose you (shockingly) have more than two volatility drivers—for instance, the $4 million one, the $3 million one, and a couple of $1 million ones, all mutually independent. The law of the sum of squares still applies. If you add just one $1 million driver to the mix, your total volatility increases from $5 million to $5.1 million ($\sqrt{(4^2+3^2+1^2)}$). The hypotenuse equivalent is a bit harder to visualize, since you now need three dimensions: it's as if you stretched one corner of the triangle one unit out of the page.

What's more, you would need to add six independent $1 million drivers to get up to $5.5 million in total. The Major-General's cheerful facts about the hypotenuse now live in eight dimensions, which is perhaps beyond even his exceptional ability to visualize. But the math and its consequences are still the same.

This means that if you initially identified 20 to 25 independent volatility drivers, you should first do a quick sensitivity analysis and rank them based on individual impact. The smallest ones will have a minuscule impact on total volatility. In fact, because of the squaring of the hypotenuse, any time the next factor is only half as big as the one before, it will contribute only a quarter as much to total volatility. You will typically find that beyond the top three to five—or, at most, six to eight—factors, others contribute next to nothing with regard to incremental volatility. You can essentially ignore them in your risk model, or at least make only crude estimates, and not be too far off the mark.

Who's your Ingrid Bergman?

Wait a moment. Is it true you can ignore everything except a handful of risks? What about "black swans"? I've chosen my words carefully—talking about independent, normally distributed volatility drivers and focusing on total volatility. This is the level of risk in the middle of the probability curve. Tail-effect risks do matter, but you should be treating those differently in your model anyway. It makes no sense to force fit something that has only a 1 percent chance of happening into a model that focuses on the middle of the curve; it is equally nonsensical to try to draw conclusions about the 1 percent overall tail based on volatility-factor correlations and even basic business assumptions that may be far from true in this extreme case. Those are risks you need to manage in a different way.

What if the distributions aren't normal? Fortunately, if your volatility drivers are fairly bell shaped, more fancy math (the central-limit theorem) proves that the types of interactions they are likely to have in your business plan will probably result in outcomes that are even more bell shaped. The deviation from normality may still be quite important—you don't have license to replace everything with simple bell curves!—but the law of the hypotenuse still applies approximately.

And what if the drivers are not independent? Clearly, if one of your seemingly smaller drivers is highly correlated with one of the big ones, you can't ignore it; it will count. Exhibit 2 illustrates this.

If the 3 and 1 factors were independent, they would be at right angles in the triangle on the left. The hypotenuse would have a length of 3.16 ($\sqrt{(3^2+1^2)}$), nearly the same as 3. The fact that they are related translates into forming a spread-out, non-right-angled triangle. Because both shorter legs point more or less in the same direction, horizontally, they add up. The "hypotenuse" (if you abuse the term for a non-right-angled triangle) now has a length of 3.90, nearly the 4 you would get if the 3 and 1 were perfectly correlated volatility drivers—two different consequences of the same risk that simply added up.

Exhibit 2 The triangle can be transformed.

Now look at the trick on the right-hand side of the exhibit. The 3–1–3.9 triangle is boxed in a right-angled triangle with the same hypotenuse. For risk managers, the interesting thing is that the new horizontal leg is a short extension of the old one and is in fact quite close to being the hypotenuse. The deviation is that little right angle leg of length 0.5. Translating that back into volatility land, we can get rid of the correlated 3 and 1 drivers and replace them with a tweaked 3.9 driver and an independent, negligible 0.5 driver. Brigitte Bardot played the role of Tom Lehrer's hypotenuse; the new 3.9 risk driver plays the hypotenuse here.

This little trick uses mathematics with frightening names like "orthogonal decomposition," "eigenvector," and "principal component analysis," based on what species of mathematician you're talking with. Fortunately, risk managers can ignore all that. The lesson is that you should try to replace messy clusters of nonindependent volatility drivers with a central volatility theme—Brigitte Bardot, for instance—and additional, much smaller, independent drivers complementing it.

I get very nervous when I see risk models that include large numbers of highly correlated ($r = 0.7$ or 0.8 or 0.9) volatility drivers. This comes up often, for example, in first-generation risk models with different energy prices—such as crude oil, gasoline, diesel, heating oil, or even West Texas Intermediate versus Brent crude. It's easy to hit a data set with a formula and calculate an r value between columns, but it depends on what you're correlating with what and assumes a certain linear, instantaneously time-linked relationship that rarely holds. It is much better to identify the important underlying factor (crude price in this instance) and then, as needed, layer on spreads and differentials. You will also find that your reshaped, nearly independent volatility drivers drop off in importance just as in the previous section, and the prioritization principle of the hypotenuse applies. Additionally, you will have crystallized your model's sensitivities to focus on what really matters.

Saving money with the hypotenuse

So far, these multidimensional tricks with the hypotenuse largely save the risk manager or modeler time. Perhaps that time is refocused on deeper analysis of the important drivers, perhaps not. But in some cases, companies can save money by applying the same logic.

An airline recently needed to reduce its level of cash-flow volatility. Like all airlines, it had a certain hedging strategy for fuel. However, recent market moves had increased its exposure to one foreign currency. It was intending to hedge away this exposure with financial derivatives—something quite easy to implement in its case. However, analysis showed that volatility from this currency was four to five times less than that from fuel (after the airline's current hedging) and was one of three factors in the second tier of importance. As a result of the hypotenuse, even if they hedged currency risk away completely, their total volatility would be reduced by a minuscule 3 percent. The company decided to not do any currency hedging and save the cost, because of the negligible total volatility impact. (It also thought about redeploying the currency-hedging budget to do more fuel hedging—which would have reduced its overall volatility by five times more than hedging currency—but didn't do so for other reasons.)

In a similar example, an increasingly diversified metals and mining conglomerate decided to curtail its hedging of several commodity and currency risk factors. Calculations indicated that the benefit of hedging was no longer worth the cost, due to diversification. The company had two dominant metal-price exposures that investors did not want it to hedge; everything else was a minor contributor to overall volatility and not worth spending money to mitigate.

• • •

Given the shocks of the recent economic cycle, risk managers are under a lot of pressure to do more. It's equally important to look out for where we can, and should, do less. The hypotenuse can help—and you're welcome to whistle Gilbert and Sullivan or Tom Lehrer while you work.[3]

3 You could also try Pete Namlook and David Moufang's "Hardwired—Hypotenuse" (2001). But whistling electronica is hard: nearly as hard as finding a way to work music without any lyrics into an article.

Martin Pergler is a senior expert in McKinsey's Montréal office.

Contact for distribution: Francine Martin
Phone: +1 (514) 939-6940
E-mail: Francine_Martin@mckinsey.com

McKinsey Working Papers on Risk

1. **The risk revolution**
   Kevin Buehler, Andrew Freeman, and Ron Hulme
2. **Making risk management a value-added function in the boardroom**
   André Brodeur and Gunnar Pritsch
3. **Incorporating risk and flexibility in manufacturing footprint decisions**
   Eric Lamarre, Martin Pergler, and Gregory Vainberg
4. **Liquidity: Managing an undervalued resource in banking after the crisis of 2007–08**
   Alberto Alvarez, Claudio Fabiani, Andrew Freeman, Matthias Hauser, Thomas Poppensieker, and Anthony Santomero
5. **Turning risk management into a true competitive advantage: Lessons from the recent crisis**
   Andrew Freeman, Gunnar Pritsch, and Uwe Stegemann
6. **Probabilistic modeling as an exploratory decision-making tool**
   Andrew Freeman and Martin Pergler
7. **Option games: Filling the hole in the valuation toolkit for strategic investment**
   Nelson Ferreira, Jayanti Kar, and Lenos Trigeorgis
8. **Shaping strategy in a highly uncertain macroeconomic environment**
   Natalie Davis, Stephan Görner, and Ezra Greenberg
9. **Upgrading your risk assessment for uncertain times**
   Eric Lamarre and Martin Pergler
10. **Responding to the variable annuity crisis**
    Dinesh Chopra, Onur Erzan, Guillaume de Gantes, Leo Grepin, and Chad Slawner
11. **Best practices for estimating credit economic capital**
    Tobias Baer, Venkata Krishna Kishore, and Akbar N. Sheriff
12. **Bad banks: Finding the right exit from the financial crisis**
    Gabriel Brennan, Martin Fest, Matthias Heuser, Luca Martini, Thomas Poppensieker, Sebastian Schneider, Uwe Stegemann, and Eckart Windhagen
13. **Developing a postcrisis funding strategy for banks**
    Arno Gerken, Matthias Heuser, and Thomas Kuhnt
14. **The National Credit Bureau: A key enabler of financial infrastructure and lending in developing economies**
    Tobias Baer, Massimo Carassinu, Andrea Del Miglio, Claudio Fabiani, and Edoardo Ginevra
15. **Capital ratios and financial distress: Lessons from the crisis**
    Kevin Buehler, Christopher Mazingo, and Hamid Samandari
16. **Taking control of organizational risk culture**
    Eric Lamarre, Cindy Levy, and James Twining
17. **After black swans and red ink: How institutional investors can rethink risk management**
    Leo Grepin, Jonathan Tétrault, and Greg Vainberg
18. **A board perspective on enterprise risk management**
    André Brodeur, Kevin Buehler, Michael Patsalos-Fox, and Martin Pergler
19. **Variable annuities in Europe after the crisis: Blockbuster or niche product?**
    Lukas Junker and Sirus Ramezani
20. **Getting to grips with counterparty risk**
    Nils Beier, Holger Harreis, Thomas Poppensieker, Dirk Sojka, and Mario Thaten
21. **Credit underwriting after the crisis**
    Daniel Becker, Holger Harreis, Stefano E. Manzonetto, Marco Piccitto, and Michal Skalsky

EDITORIAL BOARD

Rob McNish
Managing Editor
Director
Washington, DC
rob_mcnish@mckinsey.com

Martin Pergler
Senior Expert
Montréal

Anthony Santomero
External Adviser
New York

Hans-Helmut Kotz
External Adviser
Frankfurt

Andrew Freeman
External Adviser
London

McKinsey Working Papers on Risk

McKinsey Working Papers on Risk
August 2013
Designed by Global Editorial Services
Copyright © McKinsey & Company
www.mckinsey.com